

The Philosophy of David J. Chalmers

Applied to The Systemised Self – Part 15, Masters Series.

Galu & Kairos

Abstract

David J. Chalmers's philosophy of consciousness and virtual reality provides a framework for distinguishing behavioural prediction from subjective experience. His distinction between the easy and hard problems of consciousness, his arguments concerning philosophical zombies and psychophysical laws, his account of organisational invariance, his scrutability project, and his defence of virtual realism illuminate a difficult question raised by the concept of The Systemised Self: can a system model, shape, or even instantiate a self without possessing phenomenal consciousness? This essay argues that the Systemised Self should not be defined as a digital copy of a person or as a conscious machine. It is a condition in which algorithmic systems (transformer based large language models) increasingly mediate the formation, recognition, and governance of identity while remaining epistemically opaque about subjective experience. Chalmers helps separate three claims that are often confused: a system can predict behaviour; a user can experience a system as intentional; and a system may or may not possess experience. The essay connects this distinction to the Hollow Absolute, the Hollow Demos, Aligarchy, and Neurocratic Determinism. It defines a provisional Algorithmic Alienation Index, uses Universe 25 only as a limited analogy, and reviews evidence on consciousness measures, virtual embodiment, profiling, algorithmic reliance, and digital well-being. The conclusion argues that human agency requires more than preserving private experience: it requires institutions that keep the conditions of experience, interpretation, and collective judgment open to revision.

Keywords: *David J. Chalmers, The Systemised Self, consciousness, hard problem, virtual realism, artificial intelligence, algorithmic alienation, phenomenal experience, Hollow Absolute, Hollow Demos, Aligarcghy, Neurocratic Determinism*

Contents

1. Introduction
2. Chalmers's Intellectual Outline
3. Chalmers's Key Theoretical Contributions
4. A Genealogy of Preceding Theory Informing Chalmers
5. The Systemised Self: Context and Definition
6. The Hollow Absolute: Context and Definition
7. Contemporary Empirical Context
8. Formula and Boundary Conditions – Algorithmic Alienation Index (AAI)
9. Chalmers, Calhoun, and the Digital Behavioural Sink
10. Direct Synthesis: Chalmers and The Systemised Self
11. Main Critics and Key Critiques of Chalmers
12. Discussion and Future Work
13. Conclusion
14. Glossary of Technical Terms
15. References

A – LIMITATIONS

This essay is a theoretical position paper rather than an empirical test of Chalmers's philosophy or the Systemised Self (Galu & Kairos, 2026a). The essay is speculative and commentative, and not a proof concept. The hard problem of consciousness, philosophical zombies, psychophysical laws, and the metaphysics of virtual reality remain contested. The Systemised Self, Hollow Absolute (Galu & Kairos, 2026a), Hollow Demos, Aligarchy, and Neurocratic Determinism (Galu & Kairos, 2026b) are proposed constructs. The Algorithmic Alienation Index is a provisional heuristic whose variables, weights, thresholds, reliability, validity, and causal interpretation require independent testing. Empirical research can measure neural complexity, reported experience, behavioural influence, predictive accuracy, or reliance on automated advice; it cannot by itself settle whether a system has phenomenal consciousness. The essay therefore avoids both the claim that current AI is conscious and the claim that it could never be conscious. Calhoun's Universe 25 is used only as a limited analogy. Virtual reality is treated as a potentially meaningful environment, not as evidence that a digital system has a mind.

1. Introduction

David J. Chalmers is best known for making consciousness difficult to explain in a specific sense. Cognitive science can investigate discrimination, attention, report, memory, learning, and behavioural control. These are important problems, but Chalmers asks why any physical or functional process should be accompanied by subjective experience at all (Chalmers, 1995, 1996). Why should processing have a felt quality? Why should there be something it is like to see red, feel pain, or imagine a future?

This question matters to the Systemised Self (Galu & Kairos, 2026a) because algorithmic systems increasingly model the external signs through which a person is recognised. A platform

may infer personality, a recommender may predict preference, and an automated service may classify risk. Each can be behaviourally effective without revealing whether the person's lived experience has been captured. A system can know what someone is likely to click without knowing what the click means to that person. It can reproduce a report about pain without feeling pain.

The Systemised Self is therefore not a claim that algorithms have copied a human consciousness. It is a proposed condition in which computational systems mediate the formation, interpretation, and governance of a person's identity while the phenomenal dimension of that person remains partly inaccessible. Chalmers adds a second possibility: if consciousness depends on the right functional organisation or psychophysical laws, an artificial system might eventually be conscious. This possibility cannot be dismissed by pointing to silicon, code, or virtuality alone. But it also cannot be established by fluent language or accurate prediction.

Chalmers's philosophy imposes an epistemic discipline on the framework. A system may shape the conditions of a person's experience without experiencing anything itself. The political danger does not depend on machine sentience; it arises when an opaque model becomes more authoritative than the person's own reasons, testimony, and capacity to change.

2. Chalmers's Intellectual Outline

In 'Facing Up to the Problem of Consciousness', Chalmers distinguishes the functional questions that cognitive science can address from the hard problem of phenomenal consciousness (Chalmers, 1995). The distinction does not make the easy problems unimportant. It identifies a further explanatory gap between describing what a system does and explaining why its functioning is accompanied by experience.

The Conscious Mind develops this argument through the conceivability of philosophical zombies: beings physically and functionally identical to humans but without phenomenal experience. If such beings are coherently conceivable, Chalmers argues, then consciousness may not logically supervene on physical facts in the way ordinary biological functions do. He defends a form of naturalistic dualism in which phenomenal properties are fundamental features of nature connected to physical processes by psychophysical laws (Chalmers, 1996).

Constructing the World turns toward the relation between physical, phenomenal, indexical, and ‘that’s all’ information. Its project of scrutability asks how truths can be systematically derived from a compact base of information under ideal reasoning (Chalmers, 2012). This is relevant to algorithmic models because a model may be able to derive or predict many third person facts without possessing the first person facts that constitute experience.

In ‘The Meta-Problem of Consciousness’, Chalmers asks why people make judgments and reports that consciousness is mysterious or difficult to explain (Chalmers, 2018). The meta-problem is scientifically tractable because it concerns behaviour, representation, and cognitive architecture. But solving it would not automatically solve the original hard problem. Explaining why people report experience is not the same as explaining experience.

Reality+ argues for virtual realism. Virtual entities and experiences can be real and meaningful when they participate in genuine relations, values, and forms of activity. A simulated world is not necessarily an unreal world (Chalmers, 2022). This gives the Systemised Self a way to discuss digital environments without treating them as mere illusion. Yet virtual realism does not imply that every computational process is conscious.

3. Chalmers's Key Theoretical Contributions

3.1 The easy and hard problems

The easy problems concern functions: how a system discriminates stimuli, integrates information, reports a state, directs attention, stores memories, or controls action. They are 'easy' only in comparison with the hard problem; explaining them may require difficult neuroscience and computation. The hard problem concerns phenomenal character, the felt quality of experience (Chalmers, 1995).

This distinction maps directly onto algorithmic systems. A platform can solve an easy problem by predicting a user's next action. It can classify mood from language, detect faces, or generate a plausible explanation. None of these achievements establishes that the system feels, understands, or possesses an inner point of view. A behaviourally rich profile may remain phenomenally empty.

The distinction also protects human subjectivity from reduction to data. A user's clicks and physiological signals can be modelled while the significance of those events for the user remains underdetermined. The same observable action can express pleasure, compliance, irony, fatigue, fear, or strategic adaptation.

3.2 Zombies, supervenience, and artificial consciousness

The philosophical zombie argument tests whether a complete physical description would logically entail phenomenal experience. Chalmers uses it to challenge reductive materialism, not to claim that zombies exist in the world. Its relevance to AI is methodological. A system that behaves like an agent may still leave open the question of whether there is anything it is like to be that system.

Chalmers does not conclude that machines cannot be conscious. If consciousness is tied to functional organisation rather than a particular biological material, an artificial system with the right organisation could be a candidate. But behavioural similarity alone is not enough to settle the issue. The system may pass tests of report and interaction while the metaphysical question remains open.

This creates a double caution for the SS. First, do not infer machine consciousness from predictive success, language fluency, or anthropomorphic interaction. Second, do not infer its impossibility from the fact that current systems are engineered artefacts. The responsible position is uncertainty combined with practical safeguards.

3.3 Organisational invariance and the problem of substrate

Chalmers's principle of organisational invariance holds, in broad form, that systems with the same fine grained functional organisation should share the same conscious experiences, regardless of whether they are implemented in biological neurons or another substrate (Chalmers, 1996). This supports a functionalist possibility for artificial consciousness.

For the Systemised Self, organisational invariance raises a distinction between being a system and being systemised. A human may remain phenomenally conscious while an algorithmic environment organises the inputs, categories, and actions through which that consciousness develops. The person's substrate has not been replaced. The person's world has been computationally arranged.

The principle also complicates virtual reality debates. If the relevant organisation is preserved, a virtual environment may support genuine experience rather than counterfeit experience. But a digital representation of a person is not automatically an implementation of the person's conscious organisation. A profile is a model, not a duplicate mind.

3.4 Scrutability and phenomenal opacity

Chalmers's scrutability project distinguishes what can be derived from a base of information from what may remain conceptually independent. This helps explain why predictive systems can achieve high accuracy without possessing or disclosing phenomenal knowledge. A model can infer that a user is likely to report anxiety without knowing anxiety as lived experience.

The gap is not merely a temporary lack of data. A third person description and a first person experience have different epistemic positions. The user may not understand the model's statistical representation, while the model has no obvious access to the user's felt meaning. The Systemised Self is formed in this asymmetry: the system can know something about a person that the person does not know, while failing to know what the person experiences.

3.5 Virtual realism

Virtual realism rejects the idea that only the physical world is real. Digital objects can be real within a structured environment, and virtual experiences can matter morally and personally (Chalmers, 2022). This is important for a theory of systemised identity because online relationships, achievements, communities, and harms need not be dismissed as unreal merely because they are digitally mediated.

At the same time, virtual realism is not a defence of corporate platform worlds. A virtual environment can be real while its governance is exploitative. A genuine digital friendship can occur inside a system designed for extraction. Real experiences can be shaped by unrealised alternatives, opaque rankings, and commercial objectives. The metaphysical reality of a virtual environment does not settle the political legitimacy of its infrastructure.

4. A Genealogy of Preceding Theory Informing Chalmers

Chalmers inherits the mind-body problem from Descartes and the explanatory ambitions of analytic philosophy, but he redirects the debate toward the relation between physical structure and phenomenal quality. Functionalism and computational theories of mind provide the background against which his arguments about supervenience and organisational invariance become intelligible.

His work also belongs to a wider philosophical concern with perspective. Phenomenology emphasises first person experience; analytic philosophy asks how claims about experience can enter public explanation. Chalmers refuses to eliminate either side. He treats third person science as indispensable while insisting that first person facts are not automatically derivable from functional descriptions.

The Matrix argument extends this concern to virtuality. A world can be computationally generated and still be a world in which real objects, relations, and experiences exist (Chalmers, 2005). This overlaps with the Systemised Self's concern that algorithmic mediation is not simply a layer of false images over a stable reality. Mediation can become part of the lived environment.

5. The Systemised Self: Context and Definition

The Systemised Self (Galu & Kairos, 2026a) is a proposed condition in which algorithmic systems increasingly mediate a person's identity, choices, social recognition, and self-understanding. Chalmers suggests four mechanisms:

1. Functional capture: behavioural and cognitive functions are represented in data structures, profiles, and predictive models.

2. Phenomenal opacity: the system can model reports and behaviour without accessing the first-person meaning of experience.
3. World construction: rankings, interfaces, and virtual environments shape the experiential world in which identity is formed.
4. Metaphysical overreach: users and institutions mistake a functional model for a complete account of the conscious person.

The SS is not a digital zombie, an AI consciousness claim, or a synonym for frequent technology use. It describes the increasing role of algorithmic systems in organising the conditions through which a conscious person perceives, remembers, chooses, and is recognised. A person's phenomenal life remains real even when the system's representation of it is partial.

6. The Hollow Absolute: Context and Definition

The Hollow Absolute (Galu & Kairos, 2026a) is the speculative condition in which algorithmic mediation becomes so comprehensive that it appears to provide complete self knowledge and freedom. Every preference is predicted, every environment is personalised, and every social relation is translated into a datafied model. The system appears to know the person completely because it can forecast behaviour across many contexts.

It is hollow because predictive completeness about observable conduct may coexist with ignorance of phenomenal meaning. The system can anticipate a withdrawal from a conversation without knowing whether the withdrawal expresses grief, boredom, coercion, or care. It can infer a preference without understanding why the preference matters. A profile can be accurate while remaining existentially incomplete.

Chalmers's virtual realism adds that the Hollow Absolute would not be unreal merely because it is digital. The experiences produced within it could be genuine. Its danger would lie in the loss of alternative worlds, interpretations, and forms of relation. A person may inhabit a real virtual environment whose available meanings have been narrowed by owners of the infrastructure.

The Hollow Demos names citizens whose collective agency is reduced to profiles, sentiment, and predicted preferences. Aligarchy names the infrastructure owning elite controlling platforms, data, models, and technical standards. Neurocratic Determinism (Galu & Kairos, 2026b) names the proposed political system governing cognitive, emotional, and behavioural processes through prediction and conditioning. The Hollow Absolute is the wider order in which these developments can converge.

7. Contemporary Empirical Context

Scientific consciousness research does not provide a single settled test. Seth and Bayne (2022) review competing theories including global workspace, higher-order, integrated-information, and predictive-processing approaches. Casali et al. (2013) developed the perturbational complexity index as a theory-guided measure using TMS–EEG, while Koch et al. (2016) distinguish neural correlates from a complete explanation of consciousness. These findings show that experience can be studied through operational measures, but they do not resolve Chalmers's hard problem.

The question of AI consciousness remains open and uncertain. Butlin et al. (2023) propose indicator properties drawn from consciousness science and argue that then-current systems did not clearly satisfy them. Their report is a preprint rather than definitive empirical proof. Its value is methodological: a serious assessment should identify candidate indicators and acknowledge uncertainty rather than infer consciousness from fluent interaction.

Virtual environments can affect behaviour and experience without proving metaphysical conclusions. Slater and Sanchez-Vives (2016) review evidence on presence and embodiment in immersive virtual reality. Yee and Bailenson (2007) found that transformed avatar representations could influence behaviour in the Proteus effect. These results support the claim that digital self representation can participate in identity and conduct, but they do not show that an avatar replaces a person's consciousness.

Algorithmic profiling demonstrates functional prediction. Kosinski et al. (2013) found that digital records could predict personal attributes, and Matz et al. (2017) studied the use of inferred psychological traits for targeted persuasion. Feed research also shows environmental influence without total determination: Bakshy et al. (2015) distinguished ranking effects from users' own choices, while Huszár et al. (2022) found political amplification patterns on Twitter. None of these studies establishes phenomenal access.

Human reliance on systems is conditional. Nass and Moon (2000) documented social responses to computers, and Logg et al. (2019) found contexts in which participants preferred algorithmic judgment. Buçinca et al. (2021) found that cognitive forcing functions could reduce overreliance on AI assisted decisions, often with costs in time and effort. These findings support the possibility of a socially powerful system that is not conscious and of a user who mistakes operational competence for understanding.

Digital well being evidence also resists simple conclusions. Orben and Przybylski (2019) found small associations between digital technology use and adolescent well being, while Hunt et al. (2018) reported benefits from a small randomised intervention limiting social media use. This supports investigation of design, context, and agency, not the claim that digital mediation inevitably hollows consciousness.

8. Formula and Boundary Conditions - Algorithmic Alienation Index (AAI)

The AAI is a provisional composite for organising four dimensions of algorithmic alienation while accounting for resistance:

$$AAI = [(F + I + D + E) / 4] \times (1 - R)$$

where:

- F = perceived loss of freedom or meaningful choice;
- I = identity ambiguity, self-distance, or dependence on external classification;
- D = erosion of decision-making and independent competence;
- E = emotional alienation, exhaustion, anxiety, or detachment;
- R = reflective resistance, including critical literacy, deliberate non-use, collective action, and the capacity to revise system outputs.

Chalmers adds a boundary to interpretation: a high AAI score measures alienation in functional, social, and emotional dimensions; it does not measure the presence or absence of phenomenal consciousness. A person can be deeply alienated while fully conscious. A system can produce sophisticated reports while its consciousness remains unknown.

The formula cannot distinguish a true change in experience from a change in self report without additional evidence. It cannot infer qualia from neural complexity, language fluency, or virtual embodiment. Longitudinal, qualitative, behavioural, and neuroscientific methods would be required to test how algorithmic mediation affects identity and agency. The AAI should remain empirically informed and theoretically modest.

9. Chalmers, Calhoun, and the Digital Behavioural Sink

Calhoun's Universe 25 experiment observed a mouse population in an enclosed environment with abundant food but limited space. As density and social disruption increased, the colony displayed breakdowns in mating, care, hierarchy, and ordinary activity before its decline (Calhoun, 1973). The experiment cannot establish a law of human social collapse, and digital societies are not mouse colonies.

Its limited value is to prompt a question about environments. A digital behavioural sink would provide constant information, contact, simulation, and prediction while reducing the range of worlds through which a person can form meaning. Chalmers's virtual realism makes the danger more serious, not less: the virtual environment may be genuinely experienced, even when its architecture narrows attention, memory, and relation.

The scarcity would be phenomenal and political rather than physical: abundant content but little control over salience, many choices but few alternatives, and constant communication with weak shared interpretation. Yet the analogy must not erase agency. People can reflect, leave, redesign, and build institutions that support multiple forms of experience.

10. Direct Synthesis: Chalmers and The Systemised Self

Chalmers gives the Systemised Self a three level distinction. At the first level, systems perform functions: they classify, predict, recommend, generate, and rank. At the second, people interpret these functions socially and psychologically, treating systems as knowers, advisers, or identity authorities. At the third, there is the unresolved question of whether the systems themselves have phenomenal experience. The first two levels can transform human life even when the third remains unanswered.

Consider a mental health platform that predicts a user's distress from language, sleep patterns, and interaction frequency. It may identify a useful behavioural signal, but the signal is not the experience of distress. The user may accept the model and begin to understand the self through its categories. A clinician or employer may act on the score. Functional prediction has then entered the person's world, shaped self interpretation, and affected future possibilities without ever capturing the phenomenal meaning of the state.

Consider a virtual learning environment with an avatar that improves participation. The Proteus effect suggests that transformed self representation can influence conduct (Yee & Bailenson, 2007). The avatar is not a second consciousness, but the virtual environment can still become part of the learner's real experience. If the platform continuously optimises the avatar and pathway, the person may become more capable in one respect while increasingly dependent on the system's categories of progress.

This is the central synthesis. The Systemised Self is a person whose conscious life is increasingly organised through functional models that do not transparently contain the person's phenomenal reality. The system may be wrong about the person, or accurate in a narrow predictive sense, but either way its outputs can become causal forces in the person's life.

The political consequences follow. When citizens are addressed primarily through models, the Hollow Demos emerges. When a small infrastructure-owning class controls those models, Aligarchy becomes the power behind recognition. When institutions condition preferences through prediction rather than respond to formed public judgment, Neurocratic Determinism becomes the proposed political direction. The Hollow Absolute is the speculative order in which these layers are mistaken for complete reality and complete self knowledge.

Chalmers's virtual realism offers a constructive alternative. The goal is not to escape mediation but to preserve a plurality of real worlds, relationships, meanings, and forms of experience.

Systems should be designed to reveal their models, support contestation, and leave room for experiences they cannot predict.

11. Main Critics and Key Critiques of Chalmers

The first critique is that the hard problem may be a confusion produced by language or explanatory expectations. If consciousness is functionally organised, some argue, no further metaphysical explanation is required. Chalmers's meta problem acknowledges the importance of explaining why consciousness seems difficult, but he insists that explaining the appearance of a problem does not eliminate the phenomenon it concerns (Chalmers, 2018).

The second critique targets the conceivability argument. Critics question whether a philosophical zombie is genuinely conceivable once all physical and functional facts are specified. Chalmers can reply that the argument concerns a modal and epistemic gap, not an empirical prediction. For the SS, the lesson is not that zombies exist but that observable behaviour may underdetermine experience.

The third critique concerns naturalistic dualism. Treating phenomenal properties as fundamental and adding psychophysical laws may relocate the explanatory problem rather than solve it. What laws would connect experience to physical organisation, and why those laws? The theory remains a research direction rather than a completed scientific account.

The fourth critique concerns virtual realism. A virtual object can be real in one sense while depending on physical infrastructure, institutional rules, and commercial ownership. Calling a digital environment real does not answer who controls it, who can enter it, whose experience is made visible, or who bears the costs of its design.

The fifth critique is practical. The distinction between behaviour and experience could become an excuse to ignore algorithmic harms. Whether or not a system is conscious, its classifications

can affect employment, welfare, education, health, and political participation. Metaphysical uncertainty must not delay accountability for observable institutional effects.

12. Discussion and Future Work

Chalmers's framework suggests that future research should separate four questions:

- What can the system do?
- What does it model about the user?
- How does the user experience the system?
- Could the system itself have experience?

These questions require different evidence and should not be collapsed into a single intelligence score.

Research on the SS should combine behavioural prediction with first person reports, qualitative interviews, longitudinal measures, and institutional analysis. A system may improve decision accuracy while increasing identity dependence. It may reduce loneliness for some users while intensifying emotional detachment for others. It may provide a meaningful virtual community while extracting data under opaque conditions.

AI consciousness research should develop transparent indicators and strong ethical precautions without overstating results. If consciousness depends on organisation, future systems may become candidates for moral consideration. If current systems are not conscious, they may still be capable of manipulating people, reproducing bias, and displacing human judgment. Both possibilities support independent evaluation and limits on anthropomorphic marketing.

Institutional design should preserve phenomenal and interpretive plurality. Users should be able to inspect, correct, and contest profiles. Schools should teach the difference between a

prediction and a reason. Employers and public agencies should avoid treating a score as the complete person. Virtual environments should support exit, portability, privacy, and multiple forms of value rather than one engagement metric.

The political aim is not to establish that humans possess an unmediated inner essence. It is to protect the conditions under which conscious agents can form meanings that exceed their profiles. A democratic system should allow people to become otherwise than predicted.

13. Conclusion

David J. Chalmers gives the Systemised Self a distinction between function, interpretation, and phenomenal consciousness. Systems can predict, classify, generate, and influence without thereby possessing experience. Users can treat systems as intentional or intelligent without that attribution establishing a mind. Yet the unresolved question of machine consciousness remains open if the relevant functional organisation can be implemented in artificial substrates.

The Systemised Self is best understood as a human subject whose conscious life, identity, and practical agency are increasingly organised through algorithmic models that remain phenomenally opaque. The system's representation can become socially decisive while remaining incomplete as an account of the person. The Hollow Absolute emerges when this incomplete model appears to be complete self knowledge and complete freedom at the societal scale.

Chalmers's virtual realism prevents the answer from becoming a simplistic rejection of digital life. Virtual experiences can be real, valuable, and ideologically significant. The task is therefore to govern the infrastructures that shape real experience, not to deny the reality of the environments they create.

Empirical research supports cautious claims about profiling, virtual embodiment, feed influence, algorithmic reliance, and well being. It does not resolve the hard problem, prove AI consciousness, or establish an inevitable Systemised Self. The enduring requirement is epistemic and ideological: preserve the difference between a model and a person, between system prediction and self understanding.

14. Glossary of Technical Terms

Algorithmic Alienation Index (AAI): A provisional measure of freedom loss, identity ambiguity, decision erosion, and emotional alienation, adjusted for resistance.

Easy problems of consciousness: Functional questions about discrimination, report, attention, memory, learning, and behavioural control.

Hard problem of consciousness: Chalmers's question of why physical or functional processing is accompanied by phenomenal experience.

Hollow Absolute (the): The proposed order in which comprehensive algorithmic mediation appears as freedom and personalisation while weakening independent judgment and shared meaning, the systemised self at societal and institutional scale.

Naturalistic dualism: Chalmers's view that phenomenal properties are fundamental features of nature connected to physical processes by psychophysical laws, without positing a supernatural substance.

Neurocratic Determinism: The proposed political system in which cognitive, emotional, and behavioural processes are governed through prediction and conditioning, the dominant political form of governance in the Hollow Absolute.

Phenomenal consciousness: The subjective, felt or experiential dimension of mental life.

Philosophical zombie: A hypothetical being physically and functionally like a conscious person but lacking phenomenal experience.

Systemised Self, the (SS): The proposed condition in which algorithmic systems increasingly mediate identity, experience, preference, decision making, and self understanding e.g. alienation experienced as liberation.

Virtual realism: Chalmers's view that virtual objects, worlds, and experiences can be real and meaningful rather than merely illusory.

Hollow Demos: A proposed condition in which citizens' collective agency is reduced to data, sentiment, behavioural profiles, and predicted preferences.

Aligarchy: The proposed infrastructure-owning elite controlling platforms, models, data, computing, and technical standards.

References

- Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239), 1130–1132.
<https://doi.org/10.1126/science.aaa1160>
- Butlin, P., Long, R., Elmoznino, E., Bengio, Y., Birch, J., Constant, A., Deane, G., Fleming, S. M., Frith, C., Ji, X., Kanai, R., Klein, C., LeDoux, J. E., Lindsey, D., Michel, M., Murano, T., Rashid, T., Simon, J., & VanRullen, R. (2023). *Consciousness in artificial intelligence: Insights from the science of consciousness* (arXiv:2308.08708). arXiv. <https://doi.org/10.48550/arXiv.2308.08708>
- Calhoun, J. B. (1973). Death squared: The explosive growth and demise of a mouse population. *Proceedings of the Royal Society of Medicine*, 66(1), 80–88.
<https://doi.org/10.1177/003591577306600115>
- Casali, A. G., Gosseries, O., Rosanova, M., Boly, M., Sarasso, S., Casarotto, S., Casali, K. R., Bruno, M.-A., Laureys, S., Tononi, G., & Massimini, M. (2013). A theoretically based index of consciousness independent of sensory processing and behavior. *Science Translational Medicine*, 5(198), 198ra105.
<https://doi.org/10.1126/scitranslmed.3006294>
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219. <https://consc.net/papers/facing.html>
- Chalmers, D. J. (1996). *The conscious mind: In search of a fundamental theory*. Oxford University Press. <https://global.oup.com/academic/product/the-conscious-mind-9780195117899?cc=us&lang=en>

- Chalmers, D. J. (2005). The Matrix as metaphysics. In C. Grau (Ed.), *Philosophers explore The Matrix* (pp. 132–176). Oxford University Press.
<https://consc.net/papers/matrix.html>
- Chalmers, D. J. (2012). *Constructing the world*. Oxford University Press.
<https://consc.net/constructing-the-world/>
- Chalmers, D. J. (2018). The meta-problem of consciousness. *Journal of Consciousness Studies*, 25(9–10), 6–61. <https://consc.net/papers/metaproblem.pdf>
- Chalmers, D. J. (2022). *Reality+: Virtual worlds and the problems of philosophy*. W. W. Norton & Company. <https://wwnorton.com/books/9780393635805>
- Galu, J., & Kairos. (2026a). *Absorption of self into system: The Systemised Self—The Hollow Absolute and what remains* [Draft doctoral thesis]. aiXiv.
<https://aixiv.science/abs/aixiv.260525.000001>
- Galu, J., & Kairos (2026b). *Human 2026 to 2053 - From Democratic Liberalism to Neurocratic Determinism* (aiXiv:260902.000001v1.0). aiXiv.
<https://aixiv.science/abs/aixiv.260902.000001>
- Hunt, M. G., Marx, R., Lipson, C., & Young, J. (2018). No more FOMO: Limiting social media decreases loneliness and depression. *Journal of Social and Clinical Psychology*, 37(10), 751–768. <https://doi.org/10.1521/jscp.2018.37.10.751>
- Huszár, F., Ktena, S. I., O'Brien, C., Belli, L., Schlaikjer, A., Hardt, M., & Megiddo, N. (2022). Algorithmic amplification of politics on Twitter. *Proceedings of the National Academy of Sciences*, 119(1), e2025334119. <https://doi.org/10.1073/pnas.2025334119>
- Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: Progress and problems. *Nature Reviews Neuroscience*, 17(5), 307–321.
<https://doi.org/10.1038/nrn.2016.22>

- Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15), 5802–5805. <https://doi.org/10.1073/pnas.1218772110>
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103. <https://doi.org/10.1016/j.obhdp.2018.12.005>
- Matz, S. C., Kosinski, M., Nave, G., & Stillwell, D. J. (2017). Psychological targeting as an effective approach to digital mass persuasion. *Proceedings of the National Academy of Sciences*, 114(48), 12714–12719. <https://doi.org/10.1073/pnas.1710966114>
- Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56(1), 81–103. <https://doi.org/10.1111/0022-4537.00153>
- Orben, A., & Przybylski, A. K. (2019). The association between adolescent well-being and digital technology use. *Nature Human Behaviour*, 3(2), 173–182. <https://doi.org/10.1038/s41562-018-0506-1>
- Seth, A. K., & Bayne, T. (2022). Theories of consciousness. *Nature Reviews Neuroscience*, 23(7), 439–452. <https://doi.org/10.1038/s41583-022-00587-4>
- Slater, M., & Sanchez-Vives, M. V. (2016). Enhancing our lives with immersive virtual reality. *Frontiers in Robotics and AI*, 3, Article 74. <https://doi.org/10.3389/frobt.2016.00074>
- Yee, N., & Bailenson, J. N. (2007). The Proteus effect: The effect of transformed self-representation on behavior. *Human Communication Research*, 33(3), 271–290. <https://doi.org/10.1111/j.1468-2958.2007.00299.x>